

The Linear Model under Mixed Gaussian Inputs: Designing the Transfer Matrix

John T. Flåm, Dave Zachariah, Mikko Vehkaperä and Saikat Chatterjee

Abstract—Suppose a linear model $y = Hx + n$, where inputs x, n are independent Gaussian mixtures. The problem is to design the transfer matrix H so as to minimize the mean square error (MSE) when estimating x from y . This problem has important applications, but faces at least three hurdles. Firstly, even for a fixed H , the minimum MSE (MMSE) has no analytical form. Secondly, the MMSE is generally not convex in H . Thirdly, derivatives of the MMSE w.r.t. H are hard to obtain. This paper casts the problem as a stochastic program and invokes gradient methods.

The study is motivated by two applications in signal processing. One concerns the choice of error-reducing precoders; the other deals with selection of pilot matrices for channel estimation. In either setting, our numerical results indicate improved estimation accuracy - markedly better than those obtained by optimal design based on standard linear estimators.

Some implications of the non-convexities of the MMSE are noteworthy, yet, to our knowledge, not well known. For example, there are cases in which more pilot power is detrimental for channel estimation. This paper explains why.

Index Terms—Gaussian Mixtures, minimum mean square error (MMSE), estimation

I. PROBLEM STATEMENT

Consider the following linear system

$$y = Hx + n. \quad (1)$$

Here y is a vector of observations, and x and n are mutually independent random vectors with known Gaussian Mixture (GM) distributions:

$$x \sim \sum_{k \in K} p_k \mathcal{N}(u_x^{(k)}, C_{xx}^{(k)}) \quad (2)$$

$$n \sim \sum_{l \in L} q_l \mathcal{N}(u_n^{(l)}, C_{nn}^{(l)}). \quad (3)$$

In this work, we assume that H is a transfer matrix that we are at liberty to *design*, typically under some constraints. Specifically, our objective is to design H such that x can be estimated from y with minimum mean square error (MMSE). The MMSE, for a *fixed* H , is by definition [1]

$$\begin{aligned} \text{MMSE} &\triangleq E \{ \|x - u_{x|y}\|_2^2 \} \\ &= \iint \|x - u_{x|y}\|_2^2 f(x, y) dx dy. \end{aligned} \quad (4)$$

Here, $\|\cdot\|_2$ denotes the 2-norm, $f(x, y)$ is the joint probability density function (PDF) of (x, y) ,

$$u_{x|y} \triangleq E \{ x|y \} = \int x f(x|y) dx \quad (5)$$

is the MMSE estimator, and $f(x|y)$ is the PDF of x given y . The MMSE in equation (4) depends on H both through $u_{x|y}$ and $f(x, y)$. Our objective is to solve the following optimization problem

$$\min_{H \in \mathbb{H}} \text{MMSE}, \quad (6)$$

where \mathbb{H} denotes a prescribed set of matrices that H must belong to. Solving this optimization problem is not straightforward. In particular, three hurdles stand out. Firstly, with (2) and (3) as inputs to (1), the MMSE in (4) has no analytical closed form [2]. Thus, the effect of *any* matrix H , in terms of MMSE, cannot be evaluated exactly. Secondly, the MMSE is not convex in H . Thirdly, the first and second order derivatives of the MMSE w.r.t H cannot be calculated exactly, and accurate approximations are hard to obtain. For these reasons, and in order to make progress, we cast the problem as a stochastic program and invoke the Robbins-Monro algorithm [3], [4]. Very briefly our approach goes as follows: We draw samples from x and n and use these to compute stochastic gradients of the MMSE. These feed into an iterative gradient method that involves projection.

The contributions of the paper are several:

- As always, for greater accuracy, its preferable to use gradients instead of finite difference approximations. For this reason the paper spells out a formula for exact realization of stochastic gradients. Accordingly, the Robbins-Monro algorithm comes to replace the Kiefer-Wolfowitz procedure.
- In the design phase, we exploit the known input statistics and update H based on samples of the inputs (x, n) , instead of output y . This yields a closed form stochastic gradient, and we prove that it is unbiased.
- Numerical experiments indicate that our method has far better accuracy than methods which proceed via linear estimators. The main reason is that the optimal estimator, used here, is non-linear.
- It turns out that the non-convexities of the MMSE may have practical implications that deserve being better known. Specifically, in channel estimation, it can be harmful to increase the power of the pilot signal. This paper offers an explanation.

Clearly, in many practical problems, the quantities in (1) are complex-valued. Throughout this paper, however, they will all

John T. Flåm is with the Department of Electronics and Telecommunications, NTNU-Norwegian University of Science and Technology, Trondheim, Norway. Email: flam@iet.ntnu.no.

Dave Zachariah, Mikko Vehkaperä and Saikat Chatterjee are with the School of Electrical Engineering, KTH-Royal Institute of Technology, Sweden. Emails: davez@kth.se, mikkov@kth.se, sach@kth.se.

be assumed real. For the analysis, this assumption introduces no loss of generality, as the real and imaginary parts of (1) can always be treated separately.

The paper is organized as follows. The next section outlines two applications. It also specifies the Gaussian mixtures and motivates their use. Section III illustrates the problem by means of a simple example. Section IV spells out problem (6) in full detail. Section V reviews how the Robbins-Monro method applies. Numerical results are provided in Section VI. Section VII concludes. A large part of the detailed analysis, concerning stochastic gradients, is deferred to the appendix.

II. BACKGROUND AND MOTIVATION

The above described matrix design problem appears in various applications of interest. Next we present two of these, which are of particular interest to the signal processing community. Then we will explain and motivate the GM input statistics.

A. Linear precoder design

Consider a linear system model

$$\mathbf{y} = \underbrace{\mathbf{B}\mathbf{F}}_{\mathbf{H}} \mathbf{x} + \mathbf{n}, \quad (7)$$

where \mathbf{B} is a known matrix, and \mathbf{F} is a precoder matrix to be designed such that the mean square error (MSE) when estimating \mathbf{x} from \mathbf{y} becomes as small as possible. The vector \mathbf{n} is random noise. If \mathbf{x} and \mathbf{n} are independent and GM distributed, this is a matrix design problem as described in Section I, where \mathbf{F} is the design parameter. A typical constraint is to require that $\mathbf{F}\mathbf{x}$ cannot exceed a certain average power, i.e. $E\|\mathbf{F}\mathbf{x}\|_2^2 \leq \gamma$. Together with the nature of \mathbf{B} , this determines \mathbb{H} in (6).

A linear model with known transfer matrix \mathbf{H} and GM distributed inputs is frequently assumed within speech and image processing. In these applications, the signal of interest often exhibits multi modal behavior. That feature can be reasonably explained by assuming an underlying GM distribution. The noise is often modeled as Gaussian, which is a special case of a GM. Conveniently, with (2) and (3) as inputs, the MMSE estimator in (5) has a closed analytical form for any given \mathbf{H} . Selected works exploiting this include [5]–[9]. However, none of these works study MSE reducing precoders. These have the potential to significantly improve the estimation accuracy, and should therefore be of interest.

B. Pilot signal design

Consider a multiple-input-multiple-output (MIMO) communication model

$$\mathbf{z} = \mathbf{A}\mathbf{s} + \mathbf{n}, \quad (8)$$

where \mathbf{A} is a random channel matrix that we wish to estimate with as small MSE as possible, and \mathbf{s} is a pilot signal to be designed for that purpose. As before, \mathbf{n} is random noise. In order to estimate \mathbf{A} with some confidence, we must transmit as least as many pilot vectors as there are columns in \mathbf{A} . In addition we must assume that the realization of

\mathbf{A} does not change before all pilots have been transmitted. This assumption typically holds in flat, block-fading MIMO systems [10]–[12]. With multiple transmitted pilots, model (8) can be written in matrix form as

$$\mathbf{Z} = \mathbf{A}\mathbf{S} + \mathbf{N}. \quad (9)$$

If \mathbf{A} is $m \times n$, then this model can be vectorized into (Thm. 2, Ch. 2, [13])

$$\underbrace{\text{vec}(\mathbf{Z})}_{\mathbf{y}} = \underbrace{(\mathbf{S}^T \otimes \mathbf{I}_m)}_{\mathbf{H}} \underbrace{\text{vec}(\mathbf{A})}_{\mathbf{x}} + \underbrace{\text{vec}(\mathbf{N})}_{\mathbf{n}}. \quad (10)$$

Here \mathbf{I}_m denotes the $m \times m$ identity matrix, the $\text{vec}(\cdot)$ operator stacks the columns of a matrix into a column vector, and \otimes denotes the Kronecker product. Assuming that the channel (\mathbf{x}) and noise (\mathbf{n}) are independent and GM distributed, this is again a design problem as described in Section I, where the pilot matrix \mathbf{S} is the design parameter. A natural constraint, is to impose power limitations on the transmitted pilots, i.e. $\|\mathbf{S}\|_2^2 \leq \gamma$. Together with the structure imposed by the Kronecker product, this then determines \mathbb{H} in (6).

In (10), one may either assume that \mathbf{n} is pure background noise, or that \mathbf{n} represents noise *and* interference. In the former case a Gaussian distribution may be justifiable, whereas in the latter a GM distribution may be more suitable [14], [15]. As for the channel \mathbf{x} , a GM distribution can account for multiple fading situations. This can be useful, for example if the source is assumed to transmit from multiple locations. Then, the commonly used Rice distribution is unlikely to accurately capture the channel statistics associated with all transmit locations (especially so in urban areas). In fact, in [16] it has been experimentally observed and reported that different transmit locations are indeed associated with different channel statistics. A GM distributed channel, with multiple modes, has the potential to capture this.

The assumption that a channel realization can originate from several underlying distributions is not novel. For instance, all studies assuming channels governed by a Markov Model make this assumption, see e.g. [17], [18] and the references therein. A GM is a special case of an Hidden Markov model, where subsequent observations are independent, rather than governed by a Markov process. In spite of this, to the best of our knowledge, pilot optimization for estimating channels governed by a GM distribution has not been considered in the literature.

C. Gaussian Mixture distributions

While aimed at minimizing the MSE, most optimization studies on linear precoders [19], [20] or pilot signals [10]–[12] utilize only the first and second moments of the input distributions. Commonly, the underlying motivation is that a *linear* MMSE (LMMSE) estimator is employed. The LMMSE estimator¹ only relies on first and second order statistics, which conveniently tends to simplify the associated matrix design problem. In fact, the desired matrix can often be

¹Among all estimators which are linear (affine) in the observations, the LMMSE estimator obtains the smallest MSE.

obtained as the solution of a convex optimization problem. It is known, however, that the LMMSE estimator is optimal *only* for the special case when the random signals \mathbf{x} and \mathbf{n} are both Gaussian. For all other cases, the LMMSE estimator is suboptimal.

In practice, purely Gaussian inputs are rare. In general, the input distributions may be asymmetric, heavy tailed and/or multi modal. A type of distribution that can accommodate all of these cases is the *Gaussian Mixture* (GM) distribution. In fact, a GM can in theory represent *any* distribution with arbitrary accuracy [21], [22]. Therefore, in this work, we assume that the inputs are GM distributed as in (2) and (3). Notation (2) should be read in the distributional sense, where \mathbf{x} results from a composite experiment. First, source $k \in \mathcal{K}$ is activated with probability $p_k \geq 0$, $\sum_{k \in \mathcal{K}} p_k = 1$. Second, that source generates a Gaussian signal with distribution law $\mathcal{N}(\mathbf{u}_x^{(k)}, \mathbf{C}_{xx}^{(k)})$. For any realized \mathbf{x} , however, the underlying index k is not observable. The noise \mathbf{n} emerges in an entirely similar, but independent manner. \mathcal{K} and \mathcal{L} are index sets. In theory, it suffices that these sets are countable, but in practice they must be finite. Clearly when \mathcal{K} and \mathcal{L} are singletons, one falls back to the familiar case of Gaussian inputs.

The mixture parameters, e.g. $(p_k, \mathbf{u}_x^{(k)}, \mathbf{C}_{xx}^{(k)})_{k \in \mathcal{K}}$, are rarely given a priori. Most often they must be estimated, which is generally a non-trivial task [21], [23]. A common approach is to estimate the GM parameters from training data. The expectation maximization (EM) algorithm is well suited, and much used, for that purpose [1], [24], [25]. Briefly, the algorithm relies on observations drawn from the distribution we wish to parametrize, and some initial estimate of the parameters. The observations are used to update the parameters, iteratively, until convergence to a local maximum of the likelihood function. Because the resulting GM parameters depend on the initial estimates, the algorithm can alternatively be started from multiple initial estimates. This produces multiple sets of GM parameters, and each set can be assigned probabilities based on the training data. Our starting point is that the distributions in (2) and (3) have resulted from such, or similar model fitting.

Model (1) with GM inputs (2) and (3) is quite generic, and we have indicated two signal processing applications where the matrix design problem appears. The solution to that problem is, however, essentially application independent. It should therefore be of interest to a wide audience. To the best of our knowledge, it has not been pursued in the literature.

III. AN ILLUSTRATIVE EXAMPLE

We start by studying a special instance of the matrix design problem, where the MMSE for *all* $\mathbf{H} \in \mathbb{H}$ can be plotted. In general this is not possible, but the following simple example reveals some fundamental properties of the problem. Assume that we wish to design a precoder, as in (7), where $\mathbf{B} = \mathbf{I}_2$ and \mathbf{F} is restricted to be an orthogonal matrix. Equation (7) then simplifies to $\mathbf{y} = \mathbf{F}\mathbf{x} + \mathbf{n}$, where \mathbf{F} only *rotates* \mathbf{x} . Further, let \mathbf{x} and \mathbf{n} be independent and identically GM distributed as

$$\frac{1}{2}\mathcal{N}(\alpha\mathbf{e}_x, \mathbf{I}_2) + \frac{1}{2}\mathcal{N}(-\alpha\mathbf{e}_x, \mathbf{I}_2),$$

where α is a scalar and \mathbf{e}_x is the unit vector along the x -axis. Assume initially that $\mathbf{F} = \mathbf{I}_2$, which corresponds to no rotation. In this case, Figure 1(a) illustrates the densities of $\mathbf{F}\mathbf{x}$ (full circles) and \mathbf{n} (dashed circles), when seen from above. They are identical and sit on top of each other. Now, with a




Fig. 1. (a): Densities without any rotation. (b): The effect of rotating \mathbf{x} by $\pi/2$.

precoder that rotates \mathbf{x} by $\pi/2$, the densities of $\mathbf{F}\mathbf{x}$ and \mathbf{n} will look like in Figure 1(b). The latter configuration is preferable from an estimation viewpoint. This is clear from figure 2, where the MMSE is displayed as a function of all rotation angles between 0 and 2π (with $\alpha = 2$). As can be seen, a significant gain can be obtained by rotating $\pi/2$ (or by $3\pi/2$). This gain is *not* due to a particularly favorable signal-to-noise-ratio $\text{SNR} = E \|\mathbf{F}\mathbf{x}\|_2^2 / E \|\mathbf{n}\|_2^2$; because \mathbf{F} is orthogonal, the SNR remains equal for all rotation angles. The MMSE gain is instead due to a rotation producing a signal which *tends* to be orthogonal to the noise.

The above example is a special case of the matrix design problem, where \mathbf{H} in (1) is restricted to be orthogonal. It is clear that \mathbf{H} plays a decisive role in how accurately \mathbf{x} can be estimated. An almost equally important observation, however, is that the MMSE is not convex in \mathbf{H} . Hence, in general, we cannot expect that first order optimality (zero gradient) implies a *global* minimum. When studying the channel estimation problem further, we will see an implication of this non-convexity, which is perhaps not well known: In certain cases the MMSE of the channel estimate does not decrease with




Fig. 2. MMSE versus rotation angle.

increasing pilot power. On the contrary, the MMSE may in fact increase.

In the next section we rewrite the original minimization problem into an equivalent but more compact maximization problem. Then, in Section, V we present a stochastic optimization approach which provides a solution.

IV. AN EQUIVALENT MAXIMIZATION PROBLEM

In order to propose a solution to the matrix design problem in (6), we first rewrite expression (4). Using the results of [2], it follows that for model (1), under independent GM inputs (2) and (3), and a fixed \mathbf{H} , the MMSE can be written as

$$E \{ \|\mathbf{x} - E\{\mathbf{x}|\mathbf{y}\}\|_2^2 \} = \sum_k p_k \left(\text{tr} \left(\mathbf{C}_{\mathbf{x}\mathbf{x}}^{(k)} \right) + \left\| \mathbf{u}_{\mathbf{x}}^{(k)} \right\|_2^2 \right) - \int \left\| \mathbf{u}_{\mathbf{x}|\mathbf{y}} \right\|_2^2 f(\mathbf{y}) d\mathbf{y}. \quad (11)$$

In (11), $\text{tr}(\cdot)$ denotes the trace operator and $f(\mathbf{y})$ is a (GM) probability density function

$$f(\mathbf{y}) = \sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}), \quad (12)$$

where

$$f^{(k,l)}(\mathbf{y}) = \frac{e^{-\frac{1}{2}(\mathbf{y} - \mathbf{u}_{\mathbf{y}}^{(k,l)})^T \mathbf{C}_{\mathbf{y}\mathbf{y}}^{-(k,l)} (\mathbf{y} - \mathbf{u}_{\mathbf{y}}^{(k,l)})}}{(2\pi)^{\frac{M}{2}} |\mathbf{C}_{\mathbf{y}\mathbf{y}}^{(k,l)}|^{\frac{1}{2}}}, \quad (13)$$

$$\mathbf{u}_{\mathbf{y}}^{(k,l)} = \mathbf{H}\mathbf{u}_{\mathbf{x}}^{(k)} + \mathbf{u}_{\mathbf{n}}^{(l)}, \quad (14)$$

$$\mathbf{C}_{\mathbf{y}\mathbf{y}}^{(k,l)} = \mathbf{H}\mathbf{C}_{\mathbf{x}\mathbf{x}}^{(k)}\mathbf{H}^T + \mathbf{C}_{\mathbf{n}\mathbf{n}}^{(l)}. \quad (15)$$

In (13), $(\cdot)^T$ denotes transposition, $|\cdot|$ denotes the determinant, $\mathbf{C}_{\mathbf{y}\mathbf{y}}^{-(k,l)}$ is short for $(\mathbf{C}_{\mathbf{y}\mathbf{y}}^{(k,l)})^{-1}$ and M is the length of \mathbf{y} . The MMSE estimator $\mathbf{u}_{\mathbf{x}|\mathbf{y}}$ in (11) can be written as

$$\mathbf{u}_{\mathbf{x}|\mathbf{y}} = \frac{\sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \mathbf{u}_{\mathbf{x}}^{(k,l)}}{f(\mathbf{y})}, \quad (16)$$

where

$$\mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(k,l)} = \mathbf{u}_{\mathbf{x}}^{(k)} + \mathbf{C}_{\mathbf{x}\mathbf{x}}^{(k)} \mathbf{H}^T \mathbf{C}_{\mathbf{y}\mathbf{y}}^{-(k,l)} (\mathbf{y} - \mathbf{u}_{\mathbf{y}}^{(k,l)}). \quad (17)$$

In what follows, it is convenient to define

$$G(\mathbf{H}, \mathbf{y}) \triangleq \left\| \mathbf{u}_{\mathbf{x}|\mathbf{y}} \right\|_2^2. \quad (18)$$

This notation emphasizes that the squared norm of the MMSE estimate depends on both \mathbf{H} and the observation \mathbf{y} . In (11), only the integral depends on \mathbf{H} . Exploiting this, and using (18), the minimizer of the MMSE is that \mathbf{H} which maximizes

$$\int G(\mathbf{H}, \mathbf{y}) f(\mathbf{y}) d\mathbf{y} = E[G(\mathbf{H}, \mathbf{y})] =: g(\mathbf{H}), \quad (19)$$

subject to

$$\mathbf{H} \in \mathbb{H}. \quad (20)$$

The integral in (19) cannot be evaluated analytically, even for a fixed and known \mathbf{H} [2]. Moreover, as the example in Section III illustrated, the MMSE is generally not convex in \mathbf{H} , which implies that $g(\mathbf{H})$ is generally not concave. Hence, any optimization method that merely aims at first order optimality,

does in general not produce a global maximizer for $g(\mathbf{H})$. Finally, as argued in the appendix, neither first or second order derivatives of $g(\mathbf{H})$ w.r.t. \mathbf{H} can be computed exactly, and accurate approximations of these are hard to obtain.

V. THE ROBBINS-MONRO SOLUTION

The above observations suggest that a sampling based approach is the only viable option. The problem of maximizing a non-analytical expectation $E[G(\mathbf{H}, \mathbf{y})]$, over a parameter \mathbf{H} , falls under the umbrella of stochastic optimization. In particular, for our problem, the Robbins-Monro algorithm [3], [4], can be used to move iteratively from a judiciously chosen initial matrix \mathbf{H}_0 to a local maximizer \mathbf{H}^* . The philosophy is to update the current matrix \mathbf{H} using the gradient of the MMSE. Since the gradient cannot be calculated, one instead relies on a stochastic approximation. Translated to our problem, the idea is briefly as follows. Although (19) cannot be computed analytically, it can be estimated from independent sample vectors $\{\mathbf{y}_i = \mathbf{H}\mathbf{x}_i + \mathbf{n}_i\}_{i=1}^N$, as

$$g(\mathbf{H}) \approx \frac{1}{N} \sum_i \left\| \mathbf{u}_{\mathbf{x}|\mathbf{y}_i} \right\|_2^2. \quad (21)$$

The derivative of (21) w.r.t. \mathbf{H} represents an approximation of $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$, which can be used to update the current \mathbf{H} . Each update is then projected onto the feasible set \mathbb{H} . This is the core idea of the much celebrated Robbins-Monro algorithm [3]. In our context, the algorithm can be outlined as follows.

- Let the current matrix be \mathbf{H}_r .
- Draw at random (\mathbf{x}, \mathbf{n}) and compute $\mathbf{y} = \mathbf{H}_r \mathbf{x} + \mathbf{n}$.
- Calculate the update direction as

$$\mathbf{B}_r = \frac{\partial G(\mathbf{H}_r, \mathbf{y})}{\partial \mathbf{H}_r}, \quad (22)$$

and

$$\mathbf{W}_r = \mathbf{H}_r + \epsilon_r \mathbf{B}_r, \quad (23)$$

where $\{\epsilon_r\}_{r=1}^\infty$ is an infinite sequence of step sizes satisfying $\epsilon_r > 0$, $\epsilon_r \rightarrow 0$ and $\sum_{r=1}^\infty \epsilon_r = \infty$.

- Update the matrix as

$$\mathbf{H}_{r+1} = \pi_{\mathbb{H}}(\mathbf{W}_r), \quad (24)$$

where $\pi_{\mathbb{H}}(\cdot)$ represents the projection onto the set of permissible matrices \mathbb{H} .

- Repeat all steps until convergence.

A. Remarks on the Robbins-Monro algorithm

Recall that the input statistics (2), (3) are assumed known. Therefore, in a design phase, it is reasonable to assume that the inputs (\mathbf{x}, \mathbf{n}) can be sampled to *compute* \mathbf{y} , as indicated in the second step of the algorithm. The alternative would be to sample \mathbf{y} directly, leaving the underlying \mathbf{x} and \mathbf{n} unknown. The first approach is preferred because it guarantees that (22) becomes an unbiased estimate of $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$, whereas the alternative does not. This important point is fully explained in the appendix. In general, the Robbins-Monro procedure does not *require* observing the input realizations (\mathbf{x}, \mathbf{n}) . The algorithm converges also when only outputs \mathbf{y} are available.

For this reason we write (22) in terms of \mathbf{y} , but for our implementation we will assume that the underlying inputs (\mathbf{x}, \mathbf{n}) are fully known.

Because the gradient direction $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ in (22) is central in the algorithm, its closed form expression is derived in the appendix. Observe that $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ is random because it relies on a random realization of \mathbf{y} . Specifically it is a stochastic approximation of $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$ based on a *single* observation vector \mathbf{y} . Instantaneously, it may even point in directions opposite to $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$. In order to increase the likelihood of a beneficial update, one can alternatively compute the gradient as an average based on multiple \mathbf{y} 's, as suggested in (21). Then

$$\mathbf{B}_r = \frac{1}{N} \sum_{i=1}^N \frac{\partial G(\mathbf{H}_r, \mathbf{y}_i)}{\partial \mathbf{H}_r}.$$

In our implementation of the algorithm, however, we do not do this. In fact, it was recognized by Robbins and Monro, that choosing N large is generally inefficient. The reason is that \mathbf{H}_r is only intermediate in the calculations, and as argued in the appendix, regardless of the value of N , the update direction can be chosen such that it coincides with $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$ in expectation.

The Robbins-Monro procedure does not rely on the existence of $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ at all points. If this derivative is discontinuous, one can instead use any of its sub-gradients; all of which are defined. Consequently, if the local maximum towards which the algorithm converges has a discontinuous derivative, then the algorithm will oscillate around this point. Due to the decaying step sizes, however, the oscillations will eventually become infinitesimal, and for all practical purposes, the system comes to rest.

Convergence towards a local optimum is guaranteed only as $r \rightarrow \infty$ [3], [4]. Therefore, in theory, the algorithm must run forever in order to converge. The engineering solution, which tends to work well in practice, is to terminate the algorithm when $\|\mathbf{H}_{r+1} - \mathbf{H}_r\|_2 < \gamma$, where γ is a chosen threshold, or simply after a predefined number of iterations. Still, the associated running time may be non-negligible, and therefore the Robbins-Monro procedure is best suited when the input signals are stationary.

In general, for other problems than considered here, it may happen that the functional form of $G(\mathbf{H}, \mathbf{y})$ is unknown, even when its output can be observed for any \mathbf{H} and \mathbf{y} . In this case, $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ cannot be computed. Instead one may replace it by a *finite difference approximation* [26]. In some cases, this may also be preferable even when the derivative can be computed; Especially so if computing $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ requires much effort. When finite difference approximation are used, the procedure is known as the Kiefer-Wolfowitz algorithm [4], [26]. If the derivative can be computed, however, the Kiefer-Wolfowitz algorithm is associated with more uncertainty (larger variance) than the Robbins-Monro procedure. For the interested reader, the present paper extends [27], which considers Kiefer-Wolfowitz precoding.

VI. NUMERICAL RESULTS

In this section we will study two specific examples. One is on linear precoding, the other is on pilot design for channel

estimation. In conformance with much of the literature, we will use the *normalized* MSE (NMSE) as performance measure². This is defined as

$$\text{NMSE} = \frac{E \{ \|\mathbf{x} - \mathbf{u}_{\mathbf{x}|\mathbf{y}}\|_2^2 \}}{E \{ \|\mathbf{x}\|_2^2 \}}.$$

A. Precoder design

Here we study the performance of a Robbins-Monro precoder. As in the simple example of Section III, we restrict the precoder to be orthogonal. Thus, the norm of precoded signal is equal to that of the non-precoded signal. For the current example we choose the following parameters.

- $\mathbf{B} = \mathbf{I}_2$.
- \mathbf{x} is GM distributed with parameters

$$\begin{aligned} p_k &= 1/4, \text{ for } k = 1 \dots 4 \\ \mathbf{u}_{\mathbf{x}}^{(1)} &= \begin{bmatrix} -10 \\ 10 \end{bmatrix}, \mathbf{u}_{\mathbf{x}}^{(2)} = \begin{bmatrix} 10 \\ -10 \end{bmatrix}, \\ \mathbf{u}_{\mathbf{x}}^{(3)} &= \begin{bmatrix} 10 \\ 10 \end{bmatrix}, \mathbf{u}_{\mathbf{x}}^{(4)} = \begin{bmatrix} -10 \\ -10 \end{bmatrix}, \\ \mathbf{C}_{\mathbf{xx}}^{(1)} &= \mathbf{C}_{\mathbf{xx}}^{(2)} = \mathbf{C}_{\mathbf{xx}}^{(3)} = \mathbf{C}_{\mathbf{xx}}^{(4)} = \frac{1}{10} \mathbf{I}_2. \end{aligned}$$

- The noise is Gaussian and distributed as

$$\mathbf{n} \sim \mathcal{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, a \begin{bmatrix} 1 & 0 \\ 0 & 0.1 \end{bmatrix} \right). \quad (25)$$

where a is a scalar that can account for any chosen $\text{SNR} = \text{tr}(\mathbf{C}_{\mathbf{xx}}) / \text{tr}(\mathbf{C}_{\mathbf{nn}})$, and $\mathbf{C}_{\mathbf{xx}}$ and $\mathbf{C}_{\mathbf{nn}}$ are the covariance matrices of \mathbf{x} and \mathbf{n} respectively.

- We use $\mathbf{F}_0 = \mathbf{I}_2$ as the initial guess in the Robbins-Monro algorithm.
- As stopping criterion we use: $\|\mathbf{F}_{r+1} - \mathbf{F}_r\|_2 < 10^{-4}$.

Because we have assumed $\mathbf{B} = \mathbf{I}_2$, we may in the Robbins-Monro algorithm of Section V simply replace all \mathbf{H}_r with the precoder \mathbf{F}_r . For the projection in (24), we choose the *nearest orthogonal matrix*. This projection is the solution to the following optimization problem.

$$\mathbf{F}_{r+1} = \arg \min_{\mathbf{F} \in \mathbb{O}} \|\mathbf{F} - \mathbf{W}_r\|_2^2$$

where \mathbb{O} is the set of orthogonal matrices. The solution is particularly simple, and exploits the singular value decomposition:

$$\mathbf{W}_r = \mathbf{U} \mathbf{D} \mathbf{V}^T \Rightarrow \mathbf{F}_{r+1} = \mathbf{U} \mathbf{V}^T.$$

Figure 3 displays the NMSE with and without precoding, for increasing SNR levels. As can be seen, Robbins-Monro precoding provides a significant NMSE gain, especially at SNR levels between 0 and 10dB. Observe that the common approach of using the LMMSE estimator (and its corresponding precoder) is highly suboptimal at intermediate SNR levels. In fact, it is much worse than doing MMSE estimation without any precoding.

²Assuming \mathbf{x} to be a zero mean signal, the NMSE is never larger than 1 (zero dB). The reason is that the MMSE estimator, $\mathbf{u}_{\mathbf{x}|\mathbf{y}}$, will coincide with the prior mean of \mathbf{x} only when the SNR tends to zero. Hence, the prior mean is a worst case estimate of \mathbf{x} , and the NMSE describes the relative gain over the worst case estimate.




Fig. 3. The NMSE with and without precoding.

The above example indicates that our method generates a reasonable precoder, for *particular* GM input distributions. Admittedly, there exists input statistics for which the gain is much less significant. However, in all simulations we have carried out, the clear tendency is that a Robbins-Monro precoder/MMSE receiver outperforms the LMMSE precoder/LMMSE receiver at intermediate SNR levels.

B. Pilot design for channel estimation

Also for the channel estimation problem, the Robbins-Monro pilot matrix/MMSE receiver outperforms the LMMSE pilot matrix/LMMSE receiver at intermediate SNR levels. In the next example we choose parameters in order to highlight this, and one additional property. That property is a direct consequence of the non-convex nature of the MMSE. We believe it to be of interest, but not well known. The starting point is the channel estimation problem in (9), where we assume that all matrices are 2×2 . In the corresponding vectorized model

$$\underbrace{\text{vec}(\mathbf{Z})}_{\mathbf{y}} = \underbrace{(\mathbf{S}^T \otimes \mathbf{I}_2)}_{\mathbf{H}} \underbrace{\text{vec}(\mathbf{A})}_{\mathbf{x}} + \underbrace{\text{vec}(\mathbf{N})}_{\mathbf{n}}, \quad (26)$$

we assume the following parameters.

- The vectorized channel, \mathbf{x} , is distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I}_4)$.
- The vectorized noise, \mathbf{n} , is GM distributed with parameters

$$\begin{aligned} q_l &= 1/2, \text{ for } l = 1, 2, \\ \mathbf{u}_{\mathbf{n}}^{(1)} &= -\mathbf{u}_{\mathbf{n}}^{(2)} = 5 \begin{bmatrix} 1 & 1 & 1 & 1 \end{bmatrix}^T, \\ \mathbf{C}_{\mathbf{nn}}^{(1)} &= \mathbf{C}_{\mathbf{nn}}^{(2)} = \frac{1}{2} \mathbf{I}_4 \end{aligned} \quad (27)$$

- As constraint we impose that $\|\mathbf{S}\|_2^2 = \alpha$, where α is a positive scalar that can account for any chosen pilot power, and therefore also any $\text{SNR} = \|\mathbf{S}\|_2^2 / \text{tr}(\mathbf{C}_{\mathbf{nn}})$.
- Stopping criterion: $\|\mathbf{S}_{r+1} - \mathbf{S}_r\|_2 < 10^{-4}$.




Fig. 4. The NMSE as a function of pilot power.

As starting point for the Robbins-Monro algorithm we set \mathbf{S} equal to a scaled identity matrix satisfying the power constraint. During the iterations we rely on the following simple projection: If the candidate pilot matrix has power $\|\mathbf{S}\|_2^2 = \gamma$, then $\mathbf{S} \rightarrow \sqrt{\frac{\alpha}{\gamma}} \mathbf{S}$. Thus, if the pilot matrix does not use the entire power budget, the magnitude of all its elements are increased. Similarly, if pilot matrix has too large power, the magnitude of its elements are decreased.

Figure 4 shows the estimation accuracy for increasing SNR (increasing values of α). It can be seen that our method outperforms the commonly used LMMSE channel estimator/LMMSE pilot matrix at intermediate SNRs. In fact, for the same range of SNRs, the latter is much less than transmitting a scaled identity pilot matrix and using the MMSE estimator. As the SNR increases, however, it is known that the LMMSE estimator becomes optimal [2]. The performance gap between our approach and the LMMSE estimator at high SNR indicates that a scaled identity pilot matrix is a local optimum that the Robbins-Monro algorithm does not easily escape from. The most striking observation in figure 4, however, is that the channel estimates may become *worse* by increasing the pilot power! This is not an artifact of the Robbins-Monro algorithm; the same tendency is seen when a scaled identity (satisfying the same power constraint) is used as the pilot matrix.

C. Increased pilot power \neq improved channel estimates

We believe that the above phenomenon is not well known, and that it deserves to be explained. In order to visualize what happens, we will consider an example of smaller dimensions, but with similar properties as in the previous subsection. Specifically, we will assume that the unknown channel matrix \mathbf{A} is 2×1 , and that the pilot signal is just a scalar, $s = a$. Then, using (10) it follows that $\mathbf{H} = a\mathbf{I}_2$. Thus, in this setup, we do not *optimize* anything, we only try to explain the NMSE as function of different values for a . We assume the following parameters:

- $\mathbf{H} = a\mathbf{I}_2$, where a is a scalar that we can vary.




Fig. 5. The NMSE as a function of the scalar a .

- The signal (channel) \mathbf{x} is distributed as $\mathcal{N}(\mathbf{0}, \mathbf{I}_2)$.
- The noise \mathbf{n} is GM distributed with parameters

$$\begin{aligned} q_l &= 1/2, \text{ for } l = 1, 2, \\ \mathbf{u}_n^{(1)} &= -\mathbf{u}_n^{(2)} = 5 \begin{bmatrix} 1 & 1 \end{bmatrix}^T, \\ \mathbf{C}_{nn}^{(1)} &= \mathbf{C}_{nn}^{(2)} = \frac{1}{2} \mathbf{I}_2 \end{aligned} \quad (28)$$

In figure 5, the NMSE is plotted as a function of increasing values for the scalar a . We observe the same tendency as in figure 4: for increasing values of a (corresponding to increasing pilot power in figure 4), the NMSE may increase. In figure 5, we also plot the NMSE that would be obtained by a genie aided estimator [2]. Briefly, the genie aided estimator knows from which underlying Gaussian source the noise \mathbf{n} originates for each observation \mathbf{y} . Accordingly it can always produce the MMSE estimate corresponding to a purely Gaussian model. The genie aided estimator can of course not be implemented in practice, but because it is much better informed than the MMSE estimator, it provides a lower bound on the NMSE. Yet, from figure 5 we see that for $a < 3.45$ dB, the MMSE estimator is able to pin-point the correct noise component. A plausible explanation is the following. For small a , almost all realizations of $\mathbf{H}\mathbf{x} = a\mathbf{I}_2\mathbf{x}$ are close to the origin. Thus, observations \mathbf{y} tend to appear in two distinct clusters; one cluster centered at each noise component. As a consequence, the active noise component can essentially always be identified. As a grows, $\mathbf{H}\mathbf{x} = a\mathbf{I}_2\mathbf{x}$ take values in an increasingly larger area, and the cluster borders approach each other. The value $a = 3.45$ dB is the largest value for a where the clusters can still be 'perfectly' separated. This value corresponds to the local minimum in figure 5. Because we are considering 2-dimensional random vectors, we can actually visualize these clusters. The upper part of figure 6 shows how 400 independent \mathbf{y} 's form two nearby, but separable, clusters generated at $a = 3.45$ dB. When a grows beyond this level, the receiver faces a larger identification problem: it is harder to tell which noise component was active. The lower part of figure 6 shows 400 independent \mathbf{y} 's generated at $a = 7.6$ dB. This




Fig. 6. Sampled observations \mathbf{y} for $a = 3.45$ dB and $a = 7.6$ dB.

value corresponds to the local maximum in figure 5. Here the clusters largely overlap. As a continues to grow, however, the average magnitude of a noise contribution becomes so small compared to the average magnitude of $a\mathbf{I}_2\mathbf{x}$, that near perfect recovery of \mathbf{x} eventually becomes possible.

Translated to the channel estimation problem in figure 4, the interpretation is that there is a continuous range where increasing the pilot power is harmful. From figure 4, one observes that, unless one can spend an additional 15 dB (approximately) on pilot power, one will not improve from the local minimum solution.

VII. CONCLUSION

We have provided a framework for solving the matrix design problem of the linear model under Gaussian mixture statistics. The study is motivated by two applications in signal processing. One concerns the choice of error-reducing precoders; the other deals with selection of pilot matrices for channel estimation. In either setting we use the Robbins-Monro procedure to arrive at a solution. Our numerical results indicate improved estimation accuracy at intermediate SNR levels; markedly better than those obtained by optimal design based on the LMMSE estimator.

Although the Robbins-Monro algorithm in theory only converges asymptotically, in practice we see that a hard stopping criterion may work well. The algorithm is still computationally demanding, and therefore best suited under stationary or near stationary settings.

We have explored an interesting implication of the non-convexity of the MMSE; namely a case where spending more pilot power gives worse channel estimates. This phenomenon is not linked to the stochastic optimization procedure. It can be observed without optimizing \mathbf{H} at all, and we have offered a plausible explanation.

VIII. ACKNOWLEDGMENTS

John T. Flåm is supported by the Research Council of Norway under the NORDITE/VERDIKT program, Project CROPS2 (Grant 181530/S10).

APPENDIX

This appendix derives a closed form expression for the gradient direction $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ in (22), where $G(\mathbf{H}, \mathbf{y})$ is defined through (16)-(18). To that end, it is worth observing that when optimizing \mathbf{H} it is beneficial if the designer can draw samples directly from the inputs \mathbf{x} and \mathbf{n} , and not only the output \mathbf{y} . In order to see why, assume in what follows that the order of derivation and integration can be interchanged such that

$$\begin{aligned} \frac{\partial g(\mathbf{H})}{\partial \mathbf{H}} &= \frac{\partial}{\partial \mathbf{H}} \left(\int G(\mathbf{H}, \mathbf{y}) f(\mathbf{y}) d\mathbf{y} \right) \\ &= \int \frac{\partial}{\partial \mathbf{H}} [G(\mathbf{H}, \mathbf{y}) f(\mathbf{y})] d\mathbf{y}. \end{aligned} \quad (29)$$

Now, if we can only observe outputs \mathbf{y} , we have

$$E \left(\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}} \right) = \int \frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}} f(\mathbf{y}) d\mathbf{y} \neq \frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}.$$

Hence, in this case, the update direction $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ is *not* an unbiased estimator of the gradient $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$. In contrast, assume that we can draw inputs (\mathbf{x}, \mathbf{n}) , and define

$$\|\mathbf{u}_{\mathbf{x}|\mathbf{H}\mathbf{x}+\mathbf{n}}\|_2^2 = \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n}),$$

then

$$\begin{aligned} E \left(\frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}} \right) &= \iint \frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}} f(\mathbf{x}) d\mathbf{x} f(\mathbf{n}) d\mathbf{n} \\ &= \iint \frac{\partial}{\partial \mathbf{H}} [\check{G}(\mathbf{H}\mathbf{x} + \mathbf{n}) f(\mathbf{x}) d\mathbf{x} f(\mathbf{n})] d\mathbf{n} = \frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}. \end{aligned}$$

Here, the second equality holds because $f(\mathbf{x}) d\mathbf{x} f(\mathbf{n}) d\mathbf{n}$ is independent of \mathbf{H} . Hence, $\frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}}$ is an unbiased estimator of $\frac{\partial g(\mathbf{H})}{\partial \mathbf{H}}$, which is of course desirable. Because it is beneficial to sample \mathbf{x} and \mathbf{n} , rather than just \mathbf{y} , we will assume here that the designer can do this. In practice, this implies that the optimization of \mathbf{H} is done off line, as preparation for the subsequent estimation.

In what follows, we will prove that interchanging the order of integration and derivation, as in (29) is justifiable. We will derive a closed form expression for $\frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}}$, and use this as the update direction in (22). Our strategy, however, will be to do this in the reverse order: First we compute the derivative, assuming that the change can be done, and then we show that that differentiation under the integral sign is justified. Although we assume knowledge of (\mathbf{x}, \mathbf{n}) for each observed \mathbf{y} , we will write \mathbf{y} instead of $\mathbf{H}\mathbf{x} + \mathbf{n}$, and $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ instead of $\frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}}$, simply to save space.

Using (16), $\frac{\partial G(\mathbf{H}, \mathbf{y})}{\partial \mathbf{H}}$ can be written as

$$\sum_{k,l,r,s} p_k q_l p_r q_s \frac{\partial}{\partial \mathbf{H}} \left(\frac{f^{(k,l)}(\mathbf{y}) f^{(r,s)}(\mathbf{y}) \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(k,l)T} \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)}}{\left(\sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \right)^2} \right). \quad (30)$$

In order to compute (30), we make use of the following theorem [13].

Theorem 1: For a scalar function, $\phi(\mathbf{H})$, of a matrix argument, the *differential* has the form

$$\begin{aligned} d(\phi) &= \text{tr}(\mathbf{Q}^T d(\mathbf{H})) = \text{vec}(\mathbf{Q})^T \text{vec}(d\mathbf{H}), \\ \text{where } \mathbf{Q} &= \frac{\partial \phi}{\partial \mathbf{H}}. \end{aligned}$$

In our case, we take $\phi(\mathbf{H})$ to be the expression in the large parenthesis of (30). We will identify its differential, and exploit Theorem 1 in order to obtain the derivative. For that purpose, it is convenient to define

$$f^{k,l,r,s} = f^{(k,l)}(\mathbf{y}) f^{(r,s)}(\mathbf{y}), \quad (31)$$

$$z^{k,l,r,s} = \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(k,l)T} \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)}, \quad (32)$$

$$t = \sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}). \quad (33)$$

Using these, the derivative in (30), can then be compactly written as $\frac{\partial}{\partial \mathbf{H}} \left(\frac{f^{k,l,r,s} z^{k,l,r,s}}{t^2} \right)$. Using the chain rule, the differential of this fraction is

$$\begin{aligned} d(\phi) &= d \left(\frac{f^{k,l,r,s} z^{k,l,r,s}}{t^2} \right) = - \frac{2 f^{k,l,r,s} z^{k,l,r,s} d(t)}{t^3} \\ &\quad + \frac{d(f^{k,l,r,s}) z^{k,l,r,s} + d(z^{k,l,r,s}) f^{k,l,r,s}}{t^2}. \end{aligned} \quad (34)$$

Thus, we must identify the differentials $d(f^{k,l,r,s})$, $d(z^{k,l,r,s})$ and $d(t)$, which we do in next. **Notation:** we will in the remainder of this appendix use $\langle \cdot \rangle$ to compactly denote the trace operator.

Computing $d(f^{k,l,r,s})$

$$\begin{aligned} d(f^{k,l,r,s}) &= d \left(f^{(k,l)}(\mathbf{y}) f^{(r,s)}(\mathbf{y}) \right) \\ &= d \left(f^{(k,l)}(\mathbf{y}) \right) f^{(r,s)}(\mathbf{y}) + f^{(k,l)}(\mathbf{y}) d \left(f^{(r,s)}(\mathbf{y}) \right). \end{aligned} \quad (35)$$

The differential $d(f^{(k,l)}(\mathbf{y}))$, is a differential of a Gaussian probability density function. In our case it depends on the indexes (k, l) , but in order to enhance readability, we will disregard these indexes in what follows. Thus, for now, we will use equations (12)-(17) with all indexes removed, and reincorporate the indexes when needed. In addition, we will for now disregard the constant factor $(2\pi)^{-\frac{M}{2}}$ in (13). Hence, instead of considering the differential $d(f^{(k,l)}(\mathbf{y}))$, we therefore consider $d(|\mathbf{C}_{\mathbf{y}\mathbf{y}}|^{-\frac{1}{2}} g(\mathbf{y}))$, where $g(\mathbf{y}) = e^{-\frac{1}{2}(\mathbf{y}-\mathbf{u}_{\mathbf{y}})^T \mathbf{C}_{\mathbf{y}\mathbf{y}}^{-1}(\mathbf{y}-\mathbf{u}_{\mathbf{y}})}$. This can be written as

$$\begin{aligned} &d \left(|\mathbf{C}_{\mathbf{y}\mathbf{y}}|^{-\frac{1}{2}} g(\mathbf{y}) \right) \\ &= d \left(|\mathbf{C}_{\mathbf{y}\mathbf{y}}|^{-\frac{1}{2}} \right) g(\mathbf{y}) + |\mathbf{C}_{\mathbf{y}\mathbf{y}}|^{-\frac{1}{2}} d(g(\mathbf{y})) \\ &= -\frac{g(\mathbf{y})}{2} |\mathbf{C}_{\mathbf{y}\mathbf{y}}|^{-\frac{3}{2}} d(|\mathbf{C}_{\mathbf{y}\mathbf{y}}|) \\ &\quad - \frac{g(\mathbf{y})}{2} |\mathbf{C}_{\mathbf{y}\mathbf{y}}|^{-\frac{1}{2}} d \left((\mathbf{y} - \mathbf{u}_{\mathbf{y}})^T \mathbf{C}_{\mathbf{y}\mathbf{y}}^{-1} (\mathbf{y} - \mathbf{u}_{\mathbf{y}}) \right) \end{aligned} \quad (36)$$

In the second equality we have used the chain rule, and exploited that $g(\mathbf{y})$ is an exponential function. The first

differential in (36), provided \mathbf{C}_{yy} is full rank, is (Theorem 1, ch. 8, of [13])

$$d(|\mathbf{C}_{yy}|) = |\mathbf{C}_{yy}| \langle \mathbf{C}_{yy}^{-1} d(\mathbf{C}_{yy}) \rangle \quad (37)$$

$$\begin{aligned} &= |\mathbf{C}_{yy}| \langle \mathbf{C}_{yy}^{-1} d(\mathbf{H}\mathbf{C}_{xx}\mathbf{H}^T + \mathbf{C}_{nn}) \rangle \\ &= |\mathbf{C}_{yy}| \langle \mathbf{C}_{yy}^{-1} (d(\mathbf{H})\mathbf{C}_{xx}\mathbf{H}^T + \mathbf{H}\mathbf{C}_{xx}d(\mathbf{H}^T)) \rangle \end{aligned} \quad (38)$$

$$= |\mathbf{C}_{yy}| \langle \mathbf{C}_{xx}\mathbf{H}^T\mathbf{C}_{yy}^{-1}d(\mathbf{H}) + d(\mathbf{H})\mathbf{C}_{xx}\mathbf{H}^T\mathbf{C}_{yy}^{-1} \rangle \quad (39)$$

$$= |\mathbf{C}_{yy}| \langle \mathbf{C}_{xx}\mathbf{H}^T\mathbf{C}_{yy}^{-1}d(\mathbf{H}) + \mathbf{C}_{xx}\mathbf{H}^T\mathbf{C}_{yy}^{-1}d(\mathbf{H}) \rangle \quad (40)$$

$$\begin{aligned} &= 2|\mathbf{C}_{yy}| \langle \mathbf{C}_{xx}\mathbf{H}^T\mathbf{C}_{yy}^{-1}d(\mathbf{H}) \rangle \\ &= 2|\mathbf{C}_{yy}| \langle (\mathbf{C}_{yy}^{-1}\mathbf{H}\mathbf{C}_{xx})^T d(\mathbf{H}) \rangle. \end{aligned} \quad (41)$$

In (39), we have rotated the first trace (done a cyclic permutation of the matrix product), and transposed the second trace. Because \mathbf{C}_{xx} and \mathbf{C}_{yy}^{-1} are symmetric, they are not affected by transposition. Moreover, $d(\mathbf{H}^T) = (d(\mathbf{H}))^T$. The trace operator is invariant to such rotations and transposition, and therefore these operations are justified. In (40) we have rotated the second term. Such rotations and transpositions will be frequently employed throughout. Introducing $\mathbf{w} = \mathbf{y} - \mathbf{u}_y$, the second differential of (36) can be written

$$\begin{aligned} d(\mathbf{w}^T\mathbf{C}_{yy}^{-1}\mathbf{w}) &= d(\langle \mathbf{w}^T\mathbf{C}_{yy}^{-1}\mathbf{w} \rangle) \\ &= \langle \mathbf{w}^T d(\mathbf{C}_{yy}^{-1}) \mathbf{w} \rangle + 2\langle \mathbf{w}^T\mathbf{C}_{yy}^{-1}d(\mathbf{w}) \rangle. \end{aligned} \quad (42)$$

The first term of (42) is

$$\langle \mathbf{w}^T d(\mathbf{C}_{yy}^{-1}) \mathbf{w} \rangle \quad (43)$$

$$= -\langle \mathbf{w}^T\mathbf{C}_{yy}^{-1}d(\mathbf{C}_{yy})\mathbf{C}_{yy}^{-1}\mathbf{w} \rangle \quad (44)$$

$$\begin{aligned} &= -\langle \mathbf{w}^T\mathbf{C}_{yy}^{-1}(d(\mathbf{H})\mathbf{C}_{xx}\mathbf{H}^T + \mathbf{H}\mathbf{C}_{xx}d(\mathbf{H}^T))\mathbf{C}_{yy}^{-1}\mathbf{w} \rangle \\ &= -\langle \mathbf{C}_{yy}^{-1}\mathbf{w}\mathbf{w}^T\mathbf{C}_{yy}^{-1}(d(\mathbf{H})\mathbf{C}_{xx}\mathbf{H}^T + \mathbf{H}\mathbf{C}_{xx}d(\mathbf{H}^T)) \rangle. \end{aligned} \quad (45)$$

Equation (44) results from Theorem 3, ch. 8, of [13]. Observe that $\mathbf{C}_{yy}^{-1}\mathbf{w}\mathbf{w}^T\mathbf{C}_{yy}^{-1}$ in (45) is a symmetric matrix, playing the same role as \mathbf{C}_{yy}^{-1} in (38). Therefore, we can utilize (41) and conclude that

$$\langle \mathbf{w}^T d(\mathbf{C}_{yy}^{-1}) \mathbf{w} \rangle = -2\langle (\mathbf{C}_{yy}^{-1}\mathbf{w}\mathbf{w}^T\mathbf{C}_{yy}^{-1}\mathbf{H}\mathbf{C}_{xx})^T d(\mathbf{H}) \rangle. \quad (46)$$

Recall that $\mathbf{w} = \mathbf{H}(\mathbf{x} - \mathbf{u}_x) + \mathbf{n} - \mathbf{u}_n$. The second term of (42) can therefore be written as

$$\begin{aligned} &2\langle \mathbf{w}^T\mathbf{C}_{yy}^{-1}d(\mathbf{w}) \rangle \\ &= 2\langle \mathbf{w}^T\mathbf{C}_{yy}^{-1}d(\mathbf{H})(\mathbf{x} - \mathbf{u}_x) \rangle \\ &= 2\langle (\mathbf{x} - \mathbf{u}_x)\mathbf{w}^T\mathbf{C}_{yy}^{-1}d(\mathbf{H}) \rangle \\ &= 2\langle (\mathbf{C}_{yy}^{-1}\mathbf{w}(\mathbf{x} - \mathbf{u}_x)^T)^T d(\mathbf{H}) \rangle. \end{aligned} \quad (47)$$

Using (41),(46) and (47), and inserting into (36), we find that

$$\begin{aligned} &d(|\mathbf{C}_{yy}|^{-\frac{1}{2}}g(\mathbf{y})) \\ &= -g(\mathbf{y})|\mathbf{C}_{yy}|^{-\frac{1}{2}}\langle (\mathbf{C}_{yy}^{-1}\mathbf{H}\mathbf{C}_{xx})^T d(\mathbf{H}) \rangle \\ &\quad + g(\mathbf{y})|\mathbf{C}_{yy}|^{-\frac{1}{2}}\langle (\mathbf{C}_{yy}^{-1}\mathbf{w}\mathbf{w}^T\mathbf{C}_{yy}^{-1}\mathbf{H}\mathbf{C}_{xx})^T d(\mathbf{H}) \rangle \\ &\quad - g(\mathbf{y})|\mathbf{C}_{yy}|^{-\frac{1}{2}}\langle (\mathbf{C}_{yy}^{-1}\mathbf{w}(\mathbf{x} - \mathbf{u}_x)^T)^T d(\mathbf{H}) \rangle. \end{aligned} \quad (48)$$

If we define

$$\begin{aligned} \mathbf{R}^{(k,l)} &= \mathbf{C}_{yy}^{-(k,l)}\mathbf{w}^{(k,l)}(\mathbf{x} - \mathbf{u}_x^{(k)})^T \\ &\quad + \mathbf{C}_{yy}^{-(k,l)}(\mathbf{I} - \mathbf{w}^{(k,l)}\mathbf{w}^{(k,l)T}\mathbf{C}_{yy}^{-(k,l)})\mathbf{H}\mathbf{C}_{xx}^{(k)}, \end{aligned}$$

where $\mathbf{w}^{(k,l)} = \mathbf{y} - \mathbf{u}_y^{(k,l)}$, and reincorporate the constant factor $(2\pi)^{-\frac{M}{2}}$, we now find that

$$d(f^{(k,l)}(\mathbf{y})) = -f^{(k,l)}(\mathbf{y})\langle (\mathbf{R}^{(k,l)})^T d(\mathbf{H}) \rangle. \quad (49)$$

Accordingly, (35) becomes

$$d(f^{k,l,r,s}) = \left\langle f^{k,l,r,s}(\mathbf{R}^{(k,l)} + \mathbf{R}^{(r,s)})^T d(\mathbf{H}) \right\rangle. \quad (50)$$

Computing $d(z^{k,l,r,s})$

$$\begin{aligned} d(z^{k,l,r,s}) &= d\left(\mathbf{u}_{x|y}^{(k,l)T}\mathbf{u}_{x|y}^{(r,s)}\right) \\ &= \left\langle d\left(\mathbf{u}_{x|y}^{(k,l)T}\mathbf{u}_{x|y}^{(r,s)}\right) \right\rangle \\ &= \left\langle \mathbf{u}_{x|y}^{(r,s)T}d\left(\mathbf{u}_{x|y}^{(k,l)}\right) + \mathbf{u}_{x|y}^{(k,l)T}d\left(\mathbf{u}_{x|y}^{(r,s)}\right) \right\rangle. \end{aligned} \quad (51)$$

Apart from a rearrangement of the indexes, equation (51) contains two similar terms. Hence it suffices to compute one of them. Recalling that $\mathbf{u}_{x|y}^{(k,l)}$ is defined by (17), we focus on the differential

$$\begin{aligned} &\left\langle \mathbf{u}_{x|y}^{(r,s)T}d\left(\mathbf{u}_{x|y}^{(k,l)}\right) \right\rangle \\ &= \left\langle \mathbf{u}_{x|y}^{(r,s)T}d\left(\mathbf{u}_x^{(k)} + \mathbf{C}_{xx}^{(k)}\mathbf{H}^T\mathbf{C}_{yy}^{-(k,l)}\mathbf{w}^{(k,l)}\right) \right\rangle \end{aligned}$$

$$= \left\langle \mathbf{u}_{x|y}^{(r,s)T}\mathbf{C}_{xx}^{(k)}d(\mathbf{H}^T)\mathbf{C}_{yy}^{-(k,l)}\mathbf{w}^{(k,l)} \right\rangle \quad (52)$$

$$+ \left\langle \mathbf{u}_{x|y}^{(r,s)T}\mathbf{C}_{xx}^{(k)}\mathbf{H}^T d(\mathbf{C}_{yy}^{-(k,l)})\mathbf{w}^{(k,l)} \right\rangle \quad (53)$$

$$+ \left\langle \mathbf{u}_{x|y}^{(r,s)T}\mathbf{C}_{xx}^{(k)}\mathbf{H}^T\mathbf{C}_{yy}^{-(k,l)}d(\mathbf{w}^{(k,l)}) \right\rangle. \quad (54)$$

We will resolve this term by term. The first term, (52), reads

$$\begin{aligned} &\left\langle \mathbf{u}_{x|y}^{(r,s)T}\mathbf{C}_{xx}^{(k)}d(\mathbf{H}^T)\mathbf{C}_{yy}^{-(k,l)}\mathbf{w}^{(k,l)} \right\rangle \\ &= \left\langle \left(\mathbf{C}_{yy}^{-(k,l)}\mathbf{w}^{(k,l)}\mathbf{u}_{x|y}^{(r,s)T}\mathbf{C}_{xx}^{(k)}\right)^T d(\mathbf{H}) \right\rangle. \end{aligned} \quad (55)$$

The second term, (53), can be written as

$$\begin{aligned}
& \left\langle \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)T} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T d \left(\mathbf{C}_{\mathbf{yy}}^{-(k,l)} \right) \mathbf{w}^{(k,l)} \right\rangle \\
&= - \left\langle \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)T} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T \mathbf{C}_{\mathbf{yy}}^{-(k,l)} d \left(\mathbf{C}_{\mathbf{yy}}^{(k,l)} \right) \mathbf{C}_{\mathbf{yy}}^{-(k,l)} \mathbf{w}^{(k,l)} \right\rangle \\
&= - \left\langle \underbrace{\mathbf{C}_{\mathbf{yy}}^{-(k,l)} \mathbf{w}^{(k,l)} \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)T} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T \mathbf{C}_{\mathbf{yy}}^{-(k,l)}}_{\mathbf{C}^{(k,l,r,s)}} d \left(\mathbf{C}_{\mathbf{yy}}^{(k,l)} \right) \right\rangle \\
&= - \left\langle \mathbf{C}^{(k,l,r,s)} \left(d(\mathbf{H}) \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T + \mathbf{H} \mathbf{C}_{\mathbf{xx}}^{(k)} d(\mathbf{H}^T) \right) \right\rangle \\
&= - \left\langle \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T \left(\mathbf{C}^{(k,l,r,s)} + \mathbf{C}^{(k,l,r,s)T} \right) d(\mathbf{H}) \right\rangle \\
&= - \left\langle \left(\left(\mathbf{C}^{(k,l,r,s)} + \mathbf{C}^{(k,l,r,s)T} \right) \mathbf{H} \mathbf{C}_{\mathbf{xx}}^{(k)} \right)^T d(\mathbf{H}) \right\rangle. \quad (56)
\end{aligned}$$

The third term, (54), reads

$$\begin{aligned}
& \left\langle \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)T} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T \mathbf{C}_{\mathbf{yy}}^{-(k,l)} d \left(\mathbf{w}^{(k,l)} \right) \right\rangle \\
&= \left\langle \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)T} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{H}^T \mathbf{C}_{\mathbf{yy}}^{-(k,l)} d(\mathbf{H}) \left(\mathbf{x} - \mathbf{u}_{\mathbf{x}}^{(k)} \right) \right\rangle \\
&= \left\langle \left(\mathbf{C}_{\mathbf{yy}}^{-(k,l)} \mathbf{H} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)} \left(\mathbf{x} - \mathbf{u}_{\mathbf{x}}^{(k)} \right)^T \right)^T d(\mathbf{H}) \right\rangle. \quad (57)
\end{aligned}$$

Using (55), (56) and (57) we now define

$$\begin{aligned}
\mathbf{D}^{(k,l,r,s)} &= \mathbf{C}_{\mathbf{yy}}^{-(k,l)} \mathbf{w}^{(k,l)} \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)T} \mathbf{C}_{\mathbf{xx}}^{(k)} \\
&\quad - \left(\mathbf{C}^{(k,l,r,s)} + \mathbf{C}^{(k,l,r,s)T} \right) \mathbf{H} \mathbf{C}_{\mathbf{xx}}^{(k)} \\
&\quad + \mathbf{C}_{\mathbf{yy}}^{-(k,l)} \mathbf{H} \mathbf{C}_{\mathbf{xx}}^{(k)} \mathbf{u}_{\mathbf{x}|\mathbf{y}}^{(r,s)} \left(\mathbf{x} - \mathbf{u}_{\mathbf{x}}^{(k)} \right)^T.
\end{aligned}$$

Due to its two similar terms, the differential in (51) can then be written

$$d(z^{k,l,r,s}) = \left\langle \left(\mathbf{D}^{(k,l,r,s)} + \mathbf{D}^{(r,s,k,l)} \right)^T d(\mathbf{H}) \right\rangle. \quad (58)$$

Computing $d(t)$

$$\begin{aligned}
d(t) &= d \left(\sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \right) = \sum_{k,l} p_k q_l d \left(f^{(k,l)}(\mathbf{y}) \right) \\
&= - \sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \left\langle \left(\mathbf{R}^{(k,l)} \right)^T d(\mathbf{H}) \right\rangle. \quad (59)
\end{aligned}$$

The last equation results immediately by employing (49).

A. Computing the derivative

Utilizing (50), (58) and (59), the complete differential in (34) can now be written as

$$\begin{aligned}
d(\phi) &= d \left(\frac{f^{k,l,r,s} z^{k,l,r,s}}{t^2} \right) \\
&= - \frac{\left\langle f^{k,l,r,s} \left(\mathbf{R}^{(k,l)} + \mathbf{R}^{(r,s)} \right)^T d(\mathbf{H}) \right\rangle z^{k,l,r,s}}{t^2} \\
&\quad + \frac{\left\langle \left(\mathbf{D}^{(k,l,r,s)} + \mathbf{D}^{(r,s,k,l)} \right)^T d(\mathbf{H}) \right\rangle f^{k,l,r,s}}{t^2} \\
&\quad + \frac{2 f^{k,l,r,s} z^{k,l,r,s} \sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \left\langle \left(\mathbf{R}^{(k,l)} \right)^T d(\mathbf{H}) \right\rangle}{t^3}. \quad (60)
\end{aligned}$$

In case of a precoder design problem, one makes the following substitutions: $\mathbf{H} = \mathbf{B}\mathbf{F}$ and $d(\mathbf{H}) = \mathbf{B}d(\mathbf{F})$ throughout. In case of the pilot design problem (10), \mathbf{H} must be substituted by $\mathbf{S}^T \otimes \mathbf{I}_m$. In addition, assuming that \mathbf{S} is $n \times r$, one makes use of the fact that

$$\begin{aligned}
\text{vec}(d\mathbf{H}) &= \text{vec} \left(d(\mathbf{S}^T) \otimes \mathbf{I}_m \right) \\
&= (\mathbf{I}_n \otimes \mathbf{K}_{mr} \otimes \mathbf{I}_m) (\mathbf{I}_{rn} \otimes \text{vec}(\mathbf{I}_m)) d(\text{vec}(\mathbf{S}^T)).
\end{aligned}$$

Here \mathbf{K}_{mr} is the Magnus and Neudecker commutation matrix [13]. Theorem 1 can then be easily applied to (60), and identifying the derivative in (30) is therefore now straightforward.

Finally, assume that \mathbf{H} in (60) is *not* a function of some other matrix. Compactly defining $p_k q_l p_r q_s f^{k,l,r,s} = h^{k,l,r,s}$, and observing that $\sum_{k,l,r,s} h^{k,l,r,s} = t^2$, we find from equations (30), (60), and Theorem 1 that

$$\begin{aligned}
\frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}} &= \\
&= - \frac{\sum_{k,l,r,s} h^{k,l,r,s} \left(\mathbf{R}^{(k,l)} + \mathbf{R}^{(r,s)} \right) z^{k,l,r,s}}{\sum_{k,l,r,s} h^{k,l,r,s}} \\
&\quad + \frac{\sum_{k,l,r,s} h^{k,l,r,s} \left(\mathbf{D}^{(k,l,r,s)} + \mathbf{D}^{(r,s,k,l)} \right)}{\sum_{k,l,r,s} h^{k,l,r,s}} \\
&\quad + \frac{2 \sum_{k,l,r,s} h^{k,l,r,s} z^{k,l,r,s} \sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \mathbf{R}^{(k,l)}}{\sum_{i,j} p_i q_j f^{(i,j)}(\mathbf{y}) \sum_{k,l,r,s} h^{k,l,r,s}}. \quad (61)
\end{aligned}$$

B. Interchanging the order of derivation and integration

Recall, that interchanging the order of derivation and integration, as in (29), was until now only assumed valid. It derives from Lebesgue's Dominated Convergence Theorem that such a change is valid if there exists a *dominating* function $v(\cdot)$ satisfying

$$\left\| \frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}} \right\|_2 \leq \|v(\mathbf{H}, \mathbf{x}, \mathbf{n})\|_2 \quad (62)$$

and

$$\iint \|v(\mathbf{H}, \mathbf{x}, \mathbf{n})\|_2 f(\mathbf{x}) f(\mathbf{n}) d\mathbf{x} d\mathbf{n} < \infty. \quad (63)$$

Now consider (61) and define the function

$$w_{k,l,r,s}(\mathbf{H}, \mathbf{x}, \mathbf{n}) = - \left(\mathbf{R}^{(k,l)} + \mathbf{R}^{(r,s)} \right) z^{k,l,r,s} \\ + \left(\mathbf{D}^{(k,l,r,s)} + \mathbf{D}^{(r,s,k,l)} \right) \\ + \frac{2z^{k,l,r,s} \sum_{k,l} p_k q_l f^{(k,l)}(\mathbf{y}) \mathbf{R}^{(k,l)}}{\sum_{i,j} p_i q_j f^{(i,j)}(\mathbf{y})}.$$

Observe that

$$\frac{\sum_{k,l,r,s} h^{k,l,r,s} w_{k,l,r,s}(\mathbf{H}, \mathbf{x}, \mathbf{n})}{\sum_{k,l,r,s} h^{k,l,r,s}} = \frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}}.$$

Hence $\frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}}$ is a convex combination of the $w_{k,l,r,s}(\mathbf{H}, \mathbf{x}, \mathbf{n})$'s, and therefore the function

$$v(\mathbf{H}, \mathbf{x}, \mathbf{n}) = \sum_{k,l,r,s} \|w_{k,l,r,s}(\mathbf{H}, \mathbf{x}, \mathbf{n})\|_2$$

clearly satisfies (62). We do not explicitly prove it her, but it can be verified that the integral

$$\iint \sum_{k,l,r,s} \|w_{k,l,r,s}(\mathbf{H}, \mathbf{x}, \mathbf{n})\|_2 f(\mathbf{x}) f(\mathbf{n}) d\mathbf{x} d\mathbf{n} \quad (64)$$

is bounded. Hence a dominating function exists, and the change of integration and derivation is justified.

C. First and second order derivatives of the objective function

When trying to compute

$$\iint \frac{\partial \check{G}(\mathbf{H}\mathbf{x} + \mathbf{n})}{\partial \mathbf{H}} f(\mathbf{x}) f(\mathbf{n}) d\mathbf{x} d\mathbf{n}, \quad (65)$$

the mixture densities in the denominators of (61) will not simplify by substitutions. An entirely similar argument provides the reason for why (19) cannot be computed analytically in the first place [2]. Hence, (65) cannot be computed analytically, and a closed form derivative of (19) w.r.t \mathbf{H} does not exist. Although not demonstrated here, a similar argument will hold also for the second order derivative.

REFERENCES

- [1] Steven M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*, Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1993.
- [2] John Flåm, Saikat Chatterjee, Kimmo Kansanen, and Torbjörn Ekman, "On MMSE Estimation - A Linear Model under Gaussian Mixture Statistics," *IEEE Transactions on Signal Processing*, 2012.
- [3] Herbert Robbins and Sutton Monro, "A stochastic approximation method," *The Annals of Mathematical Statistics*, vol. 22, no. 3, pp. 400–407, 1951.
- [4] H.J. Kushner and G. Yin, *Stochastic Approximation and Recursive Algorithms and Applications*, vol. 35, Springer Verlag, 2003.
- [5] D. Persson and T. Eriksson, "Mixture Model- and Least Squares-Based Packet Video Error Concealment," *Image Processing, IEEE Transactions on*, vol. 18, no. 5, pp. 1048–1054, may 2009.
- [6] A. Kundu, S. Chatterjee, and T.V. Sreenivas, "Subspace Based Speech Enhancement Using Gaussian Mixture Model," in *Interspeech 2008*, Brisbane, Australia, september 2008, pp. 395–398.
- [7] A. Kundu, S. Chatterjee, and T.V. Sreenivas, "Speech Enhancement using Intra-Frame Dependency in DCT Domain," in *16th European Signal Processing Conference (EUSIPCO 2008)*, Lausanne, Switzerland, August 25–29 2008.
- [8] A.D. Subramaniam, W.R. Gardner, and B.D. Rao, "Low-Complexity Source Coding Using Gaussian Mixture Models, Lattice Vector Quantization, and Recursive Coding with Application to Speech Spectrum Quantization," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 14, no. 2, pp. 524–532, march 2006.
- [9] G. Yu, G. Sapiro, and S. Mallat, "Solving Inverse Problems With Piecewise Linear Estimators: From Gaussian Mixture Models to Structured Sparsity," *Image Processing, IEEE Transactions on*, vol. PP, no. 99, pp. 1, 2011.
- [10] D. Katselis, E. Kofidis, and S. Theodoridis, "On training optimization for estimation of correlated mimo channels in the presence of multiuser interference," *Signal Processing, IEEE Transactions on*, vol. 56, no. 10, pp. 4892–4904, oct. 2008.
- [11] M. Biguesh and A.B. Gershman, "Training-based mimo channel estimation: a study of estimator tradeoffs and optimal training signals," *Signal Processing, IEEE Transactions on*, vol. 54, no. 3, pp. 884–893, march 2006.
- [12] E. Bjornson and B. Ottersten, "A framework for training-based estimation in arbitrarily correlated rician mimo channels with rician disturbance," *Signal Processing, IEEE Transactions on*, vol. 58, no. 3, pp. 1807–1820, march 2010.
- [13] Jan R. Magnus and Heinz Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics*, John Wiley & Sons, 2nd edition, 1999.
- [14] V. Cellini and G. Dona, "A novel joint channel and multi-user interference statistics estimator for uwb-ir based on gaussian mixture model," in *Ultra-Wideband, 2005. ICU 2005. 2005 IEEE International Conference on*, sept. 2005, pp. 655–660.
- [15] G.E. Healey and R. Kondepudy, "Radiometric CCD Camera Calibration and Noise Estimation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 16, no. 3, pp. 267–276, mar 1994.
- [16] J. Medbo, H. Asplund, J.E. Berg, and N. Jaldén, "Directional Channel Characteristics in Elevation and Azimuth at an Urban Macrocell Base Station," in *6th European Conference on Antennas and Propagation*, March 2012.
- [17] Hong Shen Wang and N. Moayeri, "Finite-state markov channel-a useful model for radio communication channels," *Vehicular Technology, IEEE Transactions on*, vol. 44, no. 1, pp. 163–171, feb 1995.
- [18] P. Sadeghi, R. Kennedy, P. Rapajic, and R. Shams, "Finite-state markov modeling of fading channels - a survey of principles and applications," *Signal Processing Magazine, IEEE*, vol. 25, no. 5, pp. 57–80, september 2008.
- [19] A. Scaglione, P. Stoica, S. Barbarossa, G.B. Giannakis, and H. Sampath, "Optimal Designs for Space-Time Linear Precoders and Decoders," *Signal Processing, IEEE Transactions on*, vol. 50, no. 5, pp. 1051–1064, may 2002.
- [20] D.H. Pham, H.D. Tuan, B.-N. Vo, and T.Q. Nguyen, "Jointly optimal precoding/postcoding for colored mimo systems," in *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, may 2006, vol. 4, p. IV.
- [21] Jonathan Q. Li and Andrew R. Barron, "Mixture Density Estimation," in *In Advances in Neural Information Processing Systems 12*, 1999, pp. 279–285, MIT Press.
- [22] H.W. Sorenson and D.L. Alspach, "Recursive Bayesian Estimation Using Gaussian Sums," *Automatica*, vol. 7, no. 4, pp. 465–479, 1971.
- [23] S. Dasgupta, "Learning Mixtures of Gaussians," in *Foundations of Computer Science, 1999. 40th Annual Symposium on*, 1999, pp. 634–644.
- [24] Zoubin Ghahramani, "Solving Inverse Problems using an EM approach to Density Estimation," in *Proceedings of the 1993 Connectionist Models Summer School*, 1993, pp. 316–323.
- [25] David J.C. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003.
- [26] J. Kiefer and J. Wolfowitz, "Stochastic Estimation of the Maximum of a Regression Function," *The Annals of Mathematical Statistics*, vol. 23, no. 3, pp. 462–466, September 1952.
- [27] J.T. Flåm, M. Vehkaperä, D. Zachariah, and E. Tsakonas, "Mean Square Error Reduction by Precoding of Mixed Gaussian Input," in *International Symposium on Information Theory and its Applications, ISITA 2012, Honolulu - Hawaii, Oct 28- Nov 01, 2012*.